Published in final edited form as:

IEEE Trans Med Imaging. 2017 January ; 36(1): 17–26. doi:10.1109/TMI.2016.2593259.

Fast Variance Prediction for Iteratively Reconstructed CT Images with Locally Quadratic Regularization

Stephen M. Schmitt [Student Member, IEEE],

Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109 USA

Mitchell M. Goodsitt, and

Department of Radiology, University of Michigan, Ann Arbor, MI 48109 USA

Jeffrey A. Fessler [Fellow, IEEE]

Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109 USA

Abstract

Predicting noise properties of iteratively reconstructed CT images is useful for analyzing reconstruction methods; for example, local noise power spectrum (NPS) predictions may be used to quantify the detectability of an image feature, to design regularization methods, or to determine dynamic tube current adjustment during a CT scan. This paper presents a method for fast prediction of reconstructed image variance and local NPS for statistical reconstruction methods using quadratic or locally quadratic regularization. Previous methods either require impractical computation times to generate an approximate map of the variance of each reconstructed voxel, or are restricted to specific CT geometries. Our method can produce a variance map of the entire image, for locally shift-invariant CT geometries with sufficiently fine angular sampling, using a computation time comparable to a single back-projection. The method requires only the projection data to be used in the reconstruction, not a reconstruction itself, and is reasonably accurate except near image edges where edge-preserving regularization behaves highly nonlinearly. We evaluate the accuracy of our method using reconstructions of both simulated CT data and real CT scans of a thorax phantom.

I. Introduction

Iterative reconstruction (IR) methods for computed tomography (CT) are receiving increased attention for their improved resolution and noise properties compared to FBP [1], [2], [3]. However, the statistical properties of IR images are difficult to compute compared to FBP. Closed-form but computationally intractable matrix expressions exist [4] for the mean and covariance matrix of the reconstructed images, so faster prediction methods are desirable. Image variance information is useful for image analysis and regularization design [5]. In addition, adjusting X-ray tube current during a scan has the potential to reduce dose [6], [7],

and with sufficiently fast variance prediction, it could be feasible to compute the proper tube current adjustment during a scan to meet a certain variance target.

One way to determine the noise level in reconstructed images would be by finding the empirical variance from an ensemble of reconstructions. This method is extremely computationally intensive, requiring numerous reconstructions, and would require unacceptable X-ray dose for *in vivo* human data. The empirical approach determines the mean and variance of all voxels simultaneously, providing complete variance maps. Prior work on variance prediction has exploited the approximate shift-invariance of projection and back-projection operations to develop DFT-based approximations for the variance map of images reconstructed from scans having locally shift-invariant system geometries with sufficiently fine angular sampling. When it is possible to assume global shift-invariance of projection and back-projection (e.g. [8], [9], for PET), the DFT of projection and back-projection can be calculated once and applied to find the variance map of an entire volume. When one can assume only *local* shift invariance rather than *global* shift invariance, then DFT-based methods are useful for theoretical analysis but require projection and back-projection for each voxel of interest to calculate the DFT [10] and are impractical in general for producing a variance map for an entire large volume. For systems having sufficient symmetries (e.g., cylindrical 3D PET [10]), one can precompute and store a collection of DFT arrays; such computation and storage appears prohibitive for axial CT and inapplicable to short-scan and helical CT. There are methods specific for 2D fan-beam [11], 3D step-and-shoot [12], 3D axial CT [13], and 3D helical CT [14] that further approximate the DFT such that the computational load of predicting variance maps for an entire volume is greatly reduced.

The main contribution of this paper is a method for predicting the variance of iteratively reconstructed CT images with locally shift-invariant scan geometries. The proposed method is significantly faster than previous methods and does not require any estimate or reconstruction of the image, only the scan geometry and the weighting data to be used in the reconstruction. We also extend our previous methods [14] to deal with general weightings in the data-fit term and to accommodate space-variant regularizers such as the uniform-resolution regularizer described in [15]. Previous methods for specific geometries [11], [13], [14] are special cases of the formulation derived here.

Section II specifies the general form of the CT reconstruction problem to which our methods apply and provides an overview of the first steps in many existing variance prediction methods. The general variance prediction methods described require the *local* frequency response (LFR) of the projection, weighting, and back-projection operator. Section III shows how we apply an approximation to the LFR, derived in the Appendix¹, to variance prediction; Section III-B provides a further simplification specific to 3DCT geometries where all of the rays have a reasonably small cone-angle, such as in axial or helical scans. Section IV validates the variance predictions by comparing with the empirical variance of images reconstructed from repeated phantom scans, in both a simulation study and a real-world study.

---

¹Supplementary material in the supplementary files / multimedia tab.

II. Background

A. Problem Domain

This paper develops methods for predicting the variance (i.e., $\text{var}(\hat{x}_j)$) of statistical image reconstructions that take the general form

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\text{argmin}} L(\mathbf{Y}; \mathbf{x}) + \alpha R(\mathbf{x}). \quad (1)$$

Here, L is a data-fit term, commonly the negative log-likelihood of the N_{obs} vectorized observations $\mathbf{Y}$ given an image vector $\mathbf{x}$ composed of N_{vox} voxel attenuations. The observation vector $\mathbf{Y}$ can represent either estimates of photon counts or of line integrals. The function $R(\mathbf{x})$ is a regularization penalty. We make the following assumptions:

1. The covariance of $\mathbf{Y}$ is diagonal, and can be estimated from the data and knowledge of the instrumentation.
2. Given an image $\mathbf{x}$, the elements of $\mathbf{Y}$ are statistically independent, and the likelihood of a particular observation Y_i is modeled in terms of the

projection $[\mathbf{Ax}]_i \triangleq \sum_{j=1}^{N_{\text{vox}}} a_{ij} x_j$, such that

$$L(\mathbf{Y}; \mathbf{x}) = \sum_{i=1}^{N_{\text{obs}}} L_i(Y_i; [\mathbf{Ax}]_i); \quad (2)$$

the $N_{\text{obs}} \times N_{\text{vox}}$ system matrix $\mathbf{A}$ has elements a_{ij} representing the projection of voxel j onto observation i . This form can account for both weighted least-squares data-fit terms, in which Y_i represents a line integral estimate and the log-likelihood is given by

$$L_i(Y_i; [\mathbf{Ax}]_i) = \frac{1}{2} w_i (Y_i - [\mathbf{Ax}]_i)^2, \quad (3)$$

and for a Poisson likelihood, in which Y_i represents a photon count, and the log-likelihood is given by [16]:

$$L_i(Y_i; [\mathbf{Ax}]_i) = \bar{Y}_i - Y_i \ln \bar{Y}_i, \quad \bar{Y}_i \triangleq b_i e^{-[\mathbf{Ax}]_i} + r_i. \quad (4)$$

It is important to note that any mismatch between L in (2) and a hypothetical “true” likelihood $\tilde{L}$ that perfectly matches the true physics of a CT system does not affect our prediction methods. Even if L is a poor approximation of $\tilde{L}$, and the resulting reconstruction $\hat{\mathbf{x}}$ is a poor image, our methods can still predict the variance of $\hat{\mathbf{x}}$, so long as L is correctly characterized. In particular, our method accurately predicts the variance

even in the usual case where there is model mismatch between the system matrix $\mathbf{A}$ and the actual scanner physics. An accurate characterization of the $\mathbf{A}$ used in a particular reconstruction algorithm is much more important to predicting the variance of that reconstruction than whether $\mathbf{A}$ accurately characterizes the physics. (Of course, the usefulness of the reconstructed image $\hat{\mathbf{x}}$ will depend on the accuracy of $\mathbf{A}$.)

3. The regularizer takes the general form

$$R(\mathbf{x}) = \sum_{d=1}^{N_C} r_d \sum_k \psi([\mathbf{C}_d \mathbf{x}]_k). \quad (5)$$

In the common case of a regularizer that penalizes first differences between neighboring voxels, d indexes the directions over which we take the differences, $\mathbf{C}_d$ is a first differencing matrix between voxels in that direction, and r_d is the relative strength of the regularizer in that direction. We assume the regularizer penalty ψ is twice-differentiable at 0, and scaled such that $\psi''(0) = 1$.

These assumptions account for many common choices of data-fit terms and regularization penalties. Under these conditions on the reconstruction, we adapt [4] to the following form for the approximate covariance of the resulting reconstruction:

$$\text{cov}(\hat{\mathbf{x}}) \approx (\mathbf{A}^T \mathbf{W} \mathbf{A} + \alpha \nabla^2 R(\check{\mathbf{x}}))^{-1} \mathbf{A}^T \hat{\mathbf{W}} \mathbf{A} \cdot (\mathbf{A}^T \mathbf{W} \mathbf{A} + \alpha \nabla^2 R(\check{\mathbf{x}}))^{-1}, \quad (6)$$

where the diagonal matrices $\mathbf{W}$ and $\hat{\mathbf{W}}$ are defined as:

$$[\mathbf{W}]_{ii} \triangleq \left. \frac{\partial^2}{\partial y^2} L_i(Y_i; y) \right|_{y=[\mathbf{A}\check{\mathbf{x}}]_i} \quad (7)$$

$$[\hat{\mathbf{W}}]_{ii} \triangleq \text{var}(Y_i) \left(\left. \frac{\partial^2}{\partial y \partial Y_i} L_i(Y_i; y) \right|_{y=[\mathbf{A}\check{\mathbf{x}}]_i} \right)^2. \quad (8)$$

In (6)–(8), $\check{\mathbf{x}}$ denotes the reconstruction using noise-free data $\check{\mathbf{Y}}$. (Our final formula will not use $\check{\mathbf{x}}$.)

For a Gaussian log-likelihood in (3) with weight w_i for observation i , the weighting matrix $\mathbf{W}$ in (7) is simply these weights as a diagonal matrix: $[\mathbf{W}]_{ii} = w_i$. The matrix $\hat{\mathbf{W}}$ has

diagonal elements $[\hat{\mathbf{W}}]_{ii} = w_i^2 \text{var}(Y_i)$; if the weights are chosen so that $w_i \approx 1/\text{var}(Y_i)$, then $\hat{\mathbf{W}} \approx \mathbf{W}$.

For the Poisson log-likelihood given in (4), $[\mathbf{W}]_{ii} = Y_i/\bar{Y}_i^2$ and $[\hat{\mathbf{W}}]_{ii} = \text{var}(Y_i)/\bar{Y}_i^2$; both are approximately $1/Y_i$ [17].

For both of these statistical models, $\bar{Y}_i$ and $\text{var}(Y_i)$ are of course unknown. However, our methods are not extremely sensitive to the particular values of $\mathbf{W}$ and $\hat{\mathbf{W}}$, so estimates of these values are acceptable [4]. Typically those estimates do not require having $\hat{\mathbf{x}}$ or $\check{\mathbf{x}}$.

We define the matrix $\mathbf{P}$ as:

$$\mathbf{P} \triangleq \sum_{d=1}^{N_C} r_d \mathbf{C}_d^T \mathbf{C}_d. \quad (9)$$

We use $\mathbf{P}$ as an approximation for the Hessian of the regularizer, $\nabla^2 R(\check{\mathbf{x}})$, in approximating (6). The actual Hessian, for regularizers of the form (5), is given by:

$$\nabla^2 R(\check{\mathbf{x}}) = \sum_{d=1}^{N_C} r_d \mathbf{C}_d^T \ddot{\Psi}_d(\check{\mathbf{x}}) \mathbf{C}_d, \quad (10)$$

where $\ddot{\Psi}_d(\check{\mathbf{x}})$ is a diagonal matrix of second derivatives of the penalty function, evaluated for an image $\check{\mathbf{x}}$:

$$[\ddot{\Psi}_d(\check{\mathbf{x}})]_{kk} \triangleq \ddot{\psi}([\mathbf{C}_d \check{\mathbf{x}}]_k). \quad (11)$$

For variance prediction, we use the approximation $\ddot{\Psi}_d(\check{\mathbf{x}}) \approx \mathbf{I}$, such that $\nabla^2 R(\check{\mathbf{x}}) \approx \mathbf{P}$. This approximation (for locally quadratic regularizers like the Huber function) is based on the idea that the majority of neighboring-voxel differences in the reconstruction $\check{\mathbf{x}}$ from noise-free projection data will be small, since the regularizer penalizes large neighboring-voxel differences. We hope that for these small differences, the second derivative of their penalties will be near 1. The utility of this assumption to our purposes of fast variance prediction is enormous. First, variance prediction using $\nabla^2 R(\check{\mathbf{x}})$ would require foreknowledge of the noiseless reconstruction $\check{\mathbf{x}}$. For non-phantom applications, $\check{\mathbf{x}}$ is unavailable. Second, even using a noisy reconstruction $\hat{\mathbf{x}}$ requires the time to compute $\hat{\mathbf{x}}$, which is much greater than the computation time of our fast methods. Using $\hat{\mathbf{x}}$ would diminish the practical utility of our methods. Any edge present breaks both the assumption that $\nabla^2 R(\check{\mathbf{x}})$ is shift-invariant, and our ability to use precomputation to accelerate our algorithm. Eqn. (9) is exactly the Hessian of R for quadratic regularization.

B. Methods

In general, the matrix $\mathbf{A}^T \mathbf{W} \mathbf{A}$ is not spatially shift-invariant. If it were shift-invariant, except for the truncation to the masked space spanned by the image vector $\mathbf{x}$, then $\mathbf{A}^T \mathbf{W} \mathbf{A}$ could be diagonalized with an n -dimensional DSFT (discrete space Fourier transform):

$$\mathbf{A}^T \mathbf{W} \mathbf{A} = \mathcal{F}^* \mathcal{D}\{H^W\} \mathcal{F},$$

where $H^W(\vec{\mathbf{v}})$ would be the n -dimensional DSFT of the impulse response of $\mathbf{A}^T \mathbf{W} \mathbf{A}$, and $\mathcal{D}$ is a “diagonalization” operator:

$$(\mathcal{D}\{H\}X)(\vec{\mathbf{v}}) = H(\vec{\mathbf{v}})X(\vec{\mathbf{v}}),$$

and $\mathcal{F}$ is the DSFT with the spatial extent limited by the image support. It is defined as:

$$(\mathcal{F}\mathbf{x})(\vec{\mathbf{v}}) = \sum_k x_k \exp(-i\vec{\mathbf{v}}^T \vec{\mathbf{n}}_k) \quad (12)$$

$$[\mathcal{F}^*X]_k = \int_{[-\frac{1}{2}, \frac{1}{2}]^n} X(\vec{\mathbf{v}}) \exp(i\vec{\mathbf{v}}^T \vec{\mathbf{n}}_k) d\vec{\mathbf{v}}, \quad (13)$$

where $\vec{\mathbf{v}}$ has units of cycles per sample. Each voxel x_j in an image $\mathbf{x}$ is centered at a spatial position denoted $\vec{\mathbf{x}}_j$, and we assume that these voxel centers are aligned on a grid such that

$$\vec{\mathbf{x}}_j = \mathbf{V} \vec{\mathbf{n}}_j + \vec{\mathbf{o}}, \quad (14)$$

where $\vec{\mathbf{n}}_j$ is a voxel index with integer coordinates, $\mathbf{V}$ is a voxel spacing matrix (e.g., $\mathbf{V} = \text{diag}(x, y, z)$ for the common choice of rectangular-cuboid voxels), and $\vec{\mathbf{o}}$ is a spatial offset.

The local impulse response (LIR) of $\mathbf{A}^T \mathbf{W} \mathbf{A}$ for the voxel j is defined by

$$\mathbf{h}_j^W \triangleq \mathbf{A}^T \mathbf{W} \mathbf{A} \mathbf{e}_j, \quad (15)$$

where $\mathbf{e}_j$ is defined as the unit vector with a single 1 at position j . This LIR can be written exactly as the impulse $\mathbf{e}_j$ operated on by a frequency-domain filter $H_j^W(\vec{\mathbf{v}})$:

$$\mathbf{h}_j^W = \mathcal{F}^* \mathcal{D}\{H_j^W\} \mathcal{F} \mathbf{e}_j, \quad (16)$$

where

$$H_j^W = \mathcal{D}\{\exp(\delta \hat{\mathbf{I}} \tilde{\mathbf{Z}} \tilde{\mathbf{e}} 2\pi \vec{\nu}^T \vec{n}_j)\} \mathcal{F} \mathbf{h}_j^W. \quad (17)$$

The diagonal term “centers” the transform at the j th voxel using the shift property of the DSFT.

We will refer to H_j^W as a *local* frequency response (LFR). In the region near voxel j , $\mathbf{A}^T \mathbf{W} \mathbf{A}$ is typically *approximately* spatially shift-invariant, leading us to an approximation

$$[\mathbf{A}^T \mathbf{W} \mathbf{A}]_{kj} \approx \mathbf{e}_k^T \mathcal{F}^* \mathcal{D}\{H_j^W\} \mathcal{F} \mathbf{e}_j, \quad (18)$$

for voxel k near voxel j , which is suggested by (15) and (16). Except at the edges of the reconstructed image, $\mathbf{P}$ can also be represented in terms of its frequency response $R(\vec{\nu})$:

$$\mathbf{P} = \mathcal{F}^* \mathcal{D}\{R\} \mathcal{F}. \quad (19)$$

With both of these matrices diagonalized, (6) approximately simplifies, locally to a voxel j , to

$$\text{cov}(\hat{\mathbf{x}}) \approx \mathcal{F}^* \mathcal{D}\{S_j\} \mathcal{F}, \quad (20)$$

where

$$S_j(\vec{\nu}) \triangleq \frac{H_j^{\hat{\mathbf{W}}}(\vec{\nu})}{(H_j^W(\vec{\nu}) + \alpha R(\vec{\nu}))^2} \quad (21)$$

is the local noise power spectrum (NPS) of the noise in the reconstruction. Representing the covariance using the NPS in (20) makes the approximation that $\mathcal{F} \mathcal{F}^*$ is an identity operation. In practice it is not exactly an identity because of the finite spatial support of the image considered by the finite sum in (12), but we follow previous work that makes similar approximations [18]. Note here the distinction between H_j^W and $H_j^{\hat{\mathbf{W}}}$; $H_j^{\hat{\mathbf{W}}}$ comes from the same derivation as (15)–(18), but with $\hat{\mathbf{W}}$ substituted for $\mathbf{W}$.

If $\mathbf{A}^T \mathbf{W} \mathbf{A}$ were shift-invariant, so that (18) were exact and not a local approximation, (21) would be the global NPS of the noise. Extracting the variance of one voxel can be done by left- and right-multiplying the covariance matrix by unit vectors:

$$\text{var}(\hat{x}_j) = \mathbf{e}_j^T \text{cov}(\hat{\mathbf{x}}) \mathbf{e}_j;$$

plugging in the approximation (20) to this expression simplifies it to an integral of the local NPS:

$$\text{var}(\hat{x}_j) \approx \int_{[-\frac{1}{2}, \frac{1}{2}]^n} S_j(\vec{\nu}) d\vec{\nu} \quad (22)$$

Prior work used (22) for variance prediction; [9] and [19] find an empirical LFR from the Fourier transform of $\mathbf{h}_j^W$, found by projecting, weighting, and back-projecting $\mathbf{e}_j$. This empirical LFR H_j^W is then numerically integrated in (22). Other work derives a closed form for H_j^W based on the specific CT geometry and projection method used for $\mathbf{A}$ and simplifies (22) for these specific realizations of H_j^W ; [11] has an approximate closed form for 2D fan-beam CT, and [12] has one for a restricted subset of 3DCT geometries. The Appendix derives an approximate closed form for H_j^W for locally shift-invariant CT geometries that we use for variance prediction.

III. Variance Prediction

A. Application to general n-dimensional CT

Our variance prediction method uses the integral (22) with the NPS expression (21). The local frequency response H_j^W is approximated using (61), found in the supplementary material, which is a factorization of H_j^W into the two terms J and E_j^W :

$$H_j^W(\vec{\nu}) \approx J(\vec{\nu}) E_j^W\left(\frac{\vec{\nu}}{\|\vec{\nu}\|}\right). \quad (23)$$

In the supplement we derive this approximate factorization by separating a continuous-operator approximation of $\mathbf{A}^T \mathbf{W} \mathbf{A}$ into a weighting-like term (that becomes E_j^W) and a continuous-operator analog of the unweighted operator $\mathbf{A}^T \mathbf{A}$. We then apply a first-order Taylor expansion that makes the projection look, local to a voxel of interest and a source position, like a single parallel-beam view with an extra magnification factor. The LFR of this single parallel-beam view can be represented as a slice through the J term; which slice that we take at a particular voxel and source position becomes part of E_j^W . The simplest version of J , for the usual rectangular-cuboid voxels, is given by (65) in the supplement as:

$$J(\vec{\nu}) = \frac{\text{sinc}^2(\nu_1) \text{sinc}^2(\nu_2) \text{sinc}^2(\nu_3)}{\|\vec{\nu}\|}.$$

This ratio contains the expected $1/\|\vec{\nu}\|$ term in the frequency response of projection followed by back-projection with no weighting. The E_j^W term contains the angle-dependent weighting and is given by (63) in the supplement as

$$E_j^W \left(\frac{\vec{v}}{\|\vec{v}\|} \right) = \sum_{\sigma \in \mathcal{B}_j \left(\frac{\vec{v}}{\|\vec{v}\|} \right)} \frac{\tilde{u}_{j,\sigma} \det(\mathbf{R}_{j,\sigma}) \det(\mathbf{V})}{\left(\frac{\partial}{\partial \sigma} \hat{\theta}_{j,\sigma}^T \mathbf{V}^{-T} \frac{\vec{v}}{\|\vec{v}\|} \right)}.$$

Here, $\tilde{u}_{j,\sigma}$ in (48) denotes the element of the statistical weights $\mathbf{W}$ corresponding to the ray from the source at position σ through the j th voxel along ray direction $\hat{\theta}_{j,\sigma}$; $\det(\mathbf{R}_{j,\sigma})$ in (40) is related to the magnification at the j th voxel of the cone-beam geometry; $\mathbf{V}$ in (14) denotes

the voxel spacing; and $\mathcal{B}_j \left(\frac{\vec{v}}{\|\vec{v}\|} \right)$ in (50) is the set of source positions σ where $\hat{\theta}_{j,\sigma}$ is perpendicular to $\mathbf{V}^{-T} \vec{v}$.

The separation (23) assumes that $\mathbf{A}^T \mathbf{W} \mathbf{A}$ can be approximated by a continuous operator, and therefore that the detector elements are closely and regularly spaced. We assume that the geometry can be approximated local to a voxel by a Taylor expansion, i.e., that the function that, for any fixed source position, maps a spatial position to its corresponding detector position, is smooth. This is reasonable for standard X-ray CT geometries built from curved or flat detectors, but might not hold for some baggage CT systems that have unusual sets of detector segments with gaps. Such gaps could preclude local shift invariance.

The utility of this factorization is that the first term, J , does not depend on the voxel location j or weighting $\mathbf{W}$; the second term, E_j^W does, but does not depend on the magnitude $\|\vec{v}\|$ of the location in frequency space and so has one fewer dimension than the argument of H_j^W .

Representing $\vec{v}$ in spherical coordinates, such that $\rho \triangleq \|\vec{v}\|$ is the spatial frequency magnitude and $\vec{\Theta} \triangleq \vec{v}/\|\vec{v}\|$ is the frequency direction, (22) becomes (see supplementary materials for details):

$$\begin{aligned} \text{var}(\hat{x}_j) &\approx \int_{\mathbb{S}^n} \int_0^{\ell_{\max}} \frac{H_j^{\hat{W}}(\rho, \vec{\Theta})}{(H_j^W(\rho, \vec{\Theta}) + \alpha R(\rho, \vec{\Theta}))^2} \rho^{n-1} d\rho d\vec{\Theta} \\ &\approx \int_{\mathbb{S}^n} \int_0^{\ell_{\max}} \frac{J(\vec{v}) E_j^{\hat{W}}(\vec{\Theta})}{(J(\vec{v}) E_j^W(\vec{\Theta}) + \alpha R(\rho, \vec{\Theta}))^2} \rho^{n-1} d\rho d\vec{\Theta} \\ &= \alpha^{-1} \left(\int_{\mathbb{S}^n} \frac{E_j^{\hat{W}}(\vec{\Theta})}{E_j^W(\vec{\Theta})} \right. \\ &\quad \left. \int_0^{\ell_{\max}} \frac{\alpha^{-1} E_j^W(\vec{\Theta}) J(\vec{v})}{(\alpha^{-1} E_j^W(\vec{\Theta}) J(\vec{v}) + R(\rho, \vec{\Theta}))^2} \rho^{n-1} d\rho d\vec{\Theta} \right) \\ &= \alpha^{-1} \int_{\mathbb{S}^n} \frac{E_j^{\hat{W}}(\vec{\Theta})}{E_j^W(\vec{\Theta})} G(\alpha^{-1} E_j^W(\vec{\Theta}), \vec{\Theta}) d\vec{\Theta}, \end{aligned} \quad (24)$$

where $n = 3$ for 3D images and $n = 2$ for 2D images. Here, the integral over $\mathbb{S}^n$ is taken over the surface of the n -dimensional sphere and $\vec{\Theta}$ represents a particular point on the surface, so

(24) is an $(n-1)$ -dimensional integral. The object-independent function $G(\gamma, \vec{\Theta})$ is defined as

$$G(\gamma, \vec{\Theta}) \triangleq \int_0^{\varrho_{\max}(\vec{\Theta})} \frac{\gamma J(\varrho, \vec{\Theta})}{(\gamma J(\varrho, \vec{\Theta}) + R(\varrho, \vec{\Theta}))^2} \varrho^{n-1} d\varrho, \quad (25)$$

where $\varrho_{\max}(\vec{\Theta})$ is the maximum extent of ϱ in $[-1/2, 1/2]^n$:

$$\varrho_{\max}(\vec{\Theta}) = \frac{1}{2\|\vec{\Theta}\|_{\infty}}.$$

In general, G cannot be computed in a closed form, but it is well-behaved and only depends on, other than its arguments γ and $\vec{\Theta}$, the regularizer, which determines $R(\vec{v})$, and the scale-invariant form of the voxel, which determines $J(\vec{v})$. The J term will be the same for all cuboids that can tile space, for example, but not the same as the J term for an image reconstructed using blob-based voxels [20].

For a particular voxel shape and regularizer, we precompute a single table of values of G and use that table to predict variance maps via (24) for multiple voxels, any regularization parameter α , any weighting $\mathbf{W}$, any voxel spacing or scan geometry.

B. Application to 3D axial and helical cone-beam CT

Section VII-F of the Appendix in the supplemental materials derives the following further approximation to the two-term factorization (23) seen in the previous section that is specific to 3DCT geometries with small cone angles:

$$H_j^W(\vec{v}) \approx J_{\text{cyl}}(\vec{v}) E_{j,\text{cyl}}^W(\Phi), \quad (26)$$

shown in (71) (in the supplementary materials). In this small cone-angle factorization, the first term $J_{\text{cyl}}(\vec{v}) = J(\vec{v}) \sec \Theta$ in (70) is again independent of voxel location j and weighting $\mathbf{W}$. The second term $E_{j,\text{cyl}}$ in (72) is:

$$E_{j,\text{cyl}}^W(\Phi) \approx \frac{D_{\text{sd}}^2 \Delta_x^3 \Delta_z}{D_{\text{so}}} \sum_{\sigma \in \mathcal{B}_j(\Phi)} \tilde{u}_{j,\sigma} \frac{\|\vec{\ell}_{j,\sigma}\|^2}{D_{2,j,\sigma}^3} |\csc(\Phi - \sigma)|,$$

where D_{sd} and D_{so} denote the distances between the source and the detector and to the isocenter respectively, $\vec{\ell}_{j,\sigma}$ in (66) denotes the ray segment from the source at σ to voxel position $\vec{x}_j$, and $D_{2,j,\sigma}$ in (67) denotes the length of the projection of that segment into the xy -plane. Unlike the E_j in the previous section, here it is a function of only one variable, the azimuthal angle Φ of $\vec{v}$ in 3D cylindrical coordinates.

These terms are derived from the factorization used in (23), applying a further approximation that takes advantage of the fact that, for a small cone angle, this spherically separable approximation is nearly cylindrically separable except close to the missing cone in frequency space around the $\mathbf{v}_3$ axis. Appendix Section VII-F shows that for the purposes of variance prediction, accuracy near this missing cone is less important.

Using the LFR approximation (26), we rearrange the variance prediction integral (22) by changing from Cartesian coordinates to cylindrical coordinates $(\rho, \Phi, \mathbf{v}_3)$, to be

$$\text{var}(\hat{x}_j) \approx \alpha^{-1} \int_0^{2\pi} \frac{E_{j,\text{cyl}}^W(\Phi)}{E_{j,\text{cyl}}^W(\Phi)} G_{\text{cyl}}(\alpha^{-1} E_{j,\text{cyl}}^W(\Phi), \Phi) d\Phi, \quad (27)$$

where we define another object-independent function G_{cyl} :

$$G_{\text{cyl}}(\gamma, \Phi) \triangleq \int_0^{\rho_{\max}(\Phi)} \int_{-\frac{1}{2}}^{\frac{1}{2}} \frac{\gamma J_{\text{cyl}}(\vec{\nu})}{-\frac{1}{2}(\gamma J_{\text{cyl}}(\vec{\nu}) + R(\vec{\nu}))^2} \rho dv_3 d\rho. \quad (28)$$

In this case, $\rho_{\max} = 1/(2 \max \{|\cos \Phi|, |\sin \Phi|\})$. Again, G_{cyl} has no closed form but is a well-behaved function of only two parameters that we precompute and tabulate. We compute this table only once for a given differencing matrix $\mathbf{C}$ and voxel shape; a particular image, weighting, system geometry, or regularization parameter α does not change the table G_{cyl} . Using the table, variance prediction via (27) simply requires looking up values of G_{cyl} and numerically integrating them in 1D. This integration can be evaluated using a coarse discretization of Φ with reasonably accurate predicted variance, especially given that the integrand is periodic and integrated over one period, a case in which numerical integration converges quickly [21]. While the method of derivation is changed, this is the form for fast variance prediction given in [14], which also reduces to the form given in [13] for quadratic regularization and an axial geometry.

For 3DCT geometries where (26) is an inaccurate approximation, such as those where a voxel under consideration has rays passing through it in directions that cover much of $\mathbb{S}^3$, one must revert to (24) for fast variance prediction.

C. Spatially Varying Regularization

In this section we consider the effect on reconstruction variance of using a spatially varying regularizer such as that defined in [15], [22] designed to produce a reconstruction with uniform spatial resolution. Each voxel has an associated factor κ_j^2 representing the ‘certainty’ of the voxel that multiplies the effect of α . This factor modulates the smoothing effect of the regularizer in otherwise less certain regions to promote uniform resolution at the cost of less uniform voxel variance. To adapt our variance prediction method to this situation, we simply define a per-voxel effective regularization parameter $\alpha_{\text{eff}} \triangleq \alpha \kappa_j^2$ and evaluate (27) with this α_{eff} .

The effect of using κ_j^2 in the regularizer is intuitive: assuming that the change in the value of G_{cyl} in (27) is small when α_{eff} is varied compared to the change in the α^{-1} multiplying the integral, the approximate variance decreases inversely with increasing certainty κ_j^2 [8].

D. Object Support Masking

Outside the support of the object there is significant approximation error because, being based on (6), our prediction method ignores the non-negativity constraint that is often used in solving the reconstruction problem (1). The empirical variance outside the object approaches zero, whereas the predicted variance is positive. We use a method similar to [23] to identify regions that are outside the support of the object and set the predicted variance in these regions to zero. Other compensation methods could be incorporated [10].

IV. Results

To evaluate our fast variance prediction approach (27), we compared it to an empirical variance map in two cases. In one case, we computed the empirical variance of reconstructions from multiple realizations of simulations of noisy projection data of an XCAT phantom. In the other case, we repeatedly scanned a physical phantom and computed the empirical variance from the reconstructions of these scans. In both cases we used the weighted least-squares data-fit term (3).

A. Simulation Data

For the simulation study, we reconstructed a $512 \times 512 \times 320$ voxel section of the XCAT phantom [24] with voxel size $x \times z = 0.9764 \times 0.625\text{mm}$ that covers an anatomical section between the neck and mid-lungs. We simulated a GE third-generation helical system geometry with a 888×64 quarter-offset detector having detector element size $1.0239 \times 1.0964\text{mm}$; the detector went through three turns with a pitch of 1, taking 2952 views. Each reconstruction used 80 iterations of an ordered-subset method [25] using 64 subsets. In the regularizer, $\mathbf{C}$ was a matrix that takes 3 first differences for each voxel, one each for the adjacent voxel in each axis. These differences were penalized by a Huber potential function:

$$\psi(x) = \begin{cases} x^2/2, & |x| \leq \delta \\ \delta|x| - \delta^2/2, & |x| > \delta, \end{cases} \quad (29)$$

where the value of δ was 10 Hounsfield units. We looked at two separate cases for regularization: one with the spatially varying regularization described in section III-C, and one without. The regularization parameter α was set empirically to 2^6 in the non-spatially-varying case, and to 2^{14} in the spatially varying case. The weighting $\mathbf{W}$ was normalized so that unattenuated rays had a weight of 1. Noise was applied to each of the simulated projections by realizing Poisson random variables with mean equal to the expected number of photons with a photon count incident on the phantom of 10^5 photons per view. The empirical standard deviation maps were produced using 89 realizations of the reconstruction in the case of non-spatially-varying regularization and 111 realizations in the case of spatially varying regularization.

Figures 1(a) (with spatially varying regularization) and 2(a) (uniform regularization) show axial, sagittal, and coronal slices of the 3D map of the empirical standard deviation from our reconstructions. Since the empirical standard deviation maps were noisy and the ground truth standard deviation is spatially slowly varying, we smoothed the empirical variance maps with a gaussian kernel with a FWHM of 3 voxels each in each direction. (Figure 7, in the supplementary material, shows a profile of the unsmoothed standard deviation.) Figures 1(b) and 2(b) show the corresponding slices through the 3D predicted standard deviation map from (27). Since standard deviation varies slowly, we computed it once per $4 \times 4 \times 4$ block and used nearest-neighbor interpolation to fill in the rest. More sophisticated interpolation could be used, but the interpolation error is minimal compared to the intrinsic approximation error of our method. Figures 1(c) and 2(c) show the ground truth XCAT phantom we used. Figures 1(d) and 2(d) show the absolute magnitude of the error of our approximated standard deviation compared to the empirical results. The gray scale in these figures is transformed to better show the dynamic range of the error. Figures 1(e) and 2(e) show the empirical and predicted standard deviation along a one-dimensional trans-axial profile through the image, behind the center of the spine in the phantom, along with the standard deviation as computed directly from (22) using (17) as the LFR (labeled ‘DFT-based’). Along this profile, dotted red lines indicate where there is an edge within one voxel. Sections near the ends of the axial FOV were omitted in all images; the empirical variance becomes extreme due to a suboptimal OS algorithm implementation that is somewhat unstable in regions where the helical sampling is poor. The OS algorithm in [26] would reduce this instability and reduce the empirical variance in the (clinically unimportant) end slices.

The computation time of our method for the entire volume using $4 \times 4 \times 4$ downsampling was 1207 CPU-seconds using one core of an Intel Core i7-860 with 16 GB of memory. The empirical reconstructions took an average of 1.71 CPU-days each using one core of an Intel X5650 processor also with 16GB of memory. Table I compares the computation time required to find the empirical variance (using 111 realizations) with the computation time required to predict the variance for the entire volume using the DFT-based method and our methods. We used the DFT-based method only to produce the one-dimensional profiles shown in Figures 1(e) and 2(e); since the computation time is large, we extrapolate to find the computation time for the entire volume for Table I.

B. Real CT scans of a thorax phantom

For our real-world dataset, we scanned a phantom 10 times with a GE Discovery CT750 HD scanner and reconstructed each of the 10 sinograms separately and produced an empirical variance map of the reconstruction. The phantom was a custom modified CIRS (Norfolk, VA) Model 003 lung nodule simulator phantom for quantitative CT [27], [28]. The geometry of the system is the same as the simulated geometry used in the previous section, with the exception of performing an axial scan using a 16-row detector and 984 views. Since we could not ensure that each scan began at the same starting angle, using multiple realizations of the same helical scan to produce an empirical variance map was not possible with our physical CT scanner. For reconstructing the axial scans, we used a projection matrix $\mathbf{A}$ that was correctly aligned to the starting angle of each scan so that each reconstruction was

aligned to the same voxel grid. We used two different tube currents (40mA, 200mA) for a low-dose and a high-dose scan, and in all cases the tube voltage was 120 kVp and the scan time was 0.5 seconds.

We reconstructed each of the 10 sinograms using statistical reconstruction methods. The size of the reconstruction was $512 \times 512 \times 32$ voxels with voxel size $x \times y \times z = 0.9764 \times 0.625 \text{mm}$, as in the simulated phantom reconstructions. Each reconstruction used 100 iterations of an ordered-subset method [25] using 64 subsets. We performed the reconstructions using two different regularizers. In the first case, the regularization used a quadratic penalty and was spatially varying using the method of Section III-C. Figure 3(c) shows three slices of a sample reconstruction using these parameters. In the second case, the penalty function used the Huber potential (29) with a threshold δ of 10 Hounsfield units and was not spatially varying. Figure 4(c) shows three slices of a sample reconstruction using these parameters. In the quadratic-penalty case, the regularization parameter α was 2^{-14} ; in the Huber-penalty case, $\alpha = 2^{28}$. In both cases, the elements of the weighting matrix $\mathbf{W}$ corresponded to the CT scanner's estimate of the inverse of the variance of each ray given the scanner-specific corrections used [29]. Given that we have several repeated scans of the same object, we computed the empirical variance of the observations $\mathbf{y}$ from this data. Using this empirical observation variance for the purposes of evaluating variance prediction would be unrealistic, since in a clinical setting we do not have this data. When using elements of the matrix $\hat{\mathbf{W}}$ for variance prediction, we estimate the observation variance from the inverse of the scanner-provided weight.

Figures 3(a) (with spatially varying, quadratic regularization) and 4(a) (uniform, Huber-penalized regularization) show axial, sagittal, and coronal slices of the 3D map of the empirical standard deviation from our real reconstructions. As in the simulated empirical standard deviation maps, the empirical maps were noisy, so we smoothed the empirical variance maps with a 2D gaussian kernel with a FWHM of 5 voxels each in each direction. (Noisy, unsmoothed maps are shown in the supplemental material.) Figures 3(b) and 4(b) show the corresponding slices through the 3D predicted standard deviation map from (27). We computed the standard deviation once per $4 \times 4 \times 1$ block and used nearest-neighbor interpolation to fill in the rest. Figures 3(d) and 4(d) show the absolute magnitude of the error of our approximated standard deviation compared to the empirical results. Figures 3(e) and 4(e) show the empirical and predicted standard deviation along a one-dimensional coronal profile through the center of the image, along with the standard deviation as computed from (22) using a DFT variant of (17) as the LFR (labeled 'DFT-based').

Table I compares the computation time required to find the empirical variance, in both the simulation study and the real-world study, with the computation time required to predict the variance for the entire volume using the DFT-based method and our methods. We used the DFT-based method only to produce the one-dimensional profiles shown in Figures 1(e), 2(e), 3(e), and 4(e); since the computation time is large, we extrapolate to find the computation time for the entire volume for Table I.

V. Discussion

The presented methods are able to predict the standard deviation of most voxels in the simulated reconstructed images within an error of one Hounsfield unit in both the spatially varying regularization case (about 85% within 1HU) and the uniform regularization case (about 95% within 1HU) in less time than empirical measurement by a factor of over 10^4 . Figure 5 shows, for the entire CT volume, the percentage of the image that had an error within a specified bound in both the spatially varying and uniform regularization cases. Whether the tradeoff for time at the expense of accuracy provided by our method is acceptable depends on the application.

Figure 3 shows very good qualitative agreement between empirical and predicted standard deviation, even away from the plane of X-ray source rotation, until the furthest end slices of the reconstruction, which have insufficient data coverage and would not be presented clinically. Inside the phantom, the region with the highest error is the center-most section. Curiously, the DFT-based method has trouble with this region as well, although not as much as our prediction. The reason for this discrepancy is unknown. One possibility is that the OS algorithm had not converged completely. Figure 4 shows good qualitative agreement between the empirical and predicted standard deviation, except near edges. Except near the edges, the standard deviation is nearly constant in both the empirical and predicted reconstructions. The predicted standard deviation has a slight variation in the profile that we can see that seems like it follows a ground-truth slight variation that exists away from edges. The prediction slightly overestimates variance in Figure 3 and underestimates Figure 4; we are not sure what causes this discrepancy. It is possible that this is a function of the regularization parameter α (which is much higher in the Huber case, to compensate for the edge-preserving effect), and that for an α somewhere between these two cases the bias crosses zero.

The comparisons in Figures 1(e), 2(e), 3(e), and 4(e) demonstrate that the majority of the error incurred by our methods occurs in the assumptions of quadratic-like regularization and local shift-invariance that ultimately lead to (22), rather than our approximations that transform (22) into the more computationally tractable (27). There is very little difference between the fast prediction and DFT-based prediction in the profile in Figure 4(e). This is reasonable considering that the error in the approximation (23) is highest when the magnitude of the frequency, $\|\mathbf{v}\|$, is large, but in these regions the regularization tends to be strongest, suppressing the error in the resulting approximate local NPS.

Since we approximate the regularizer as being locally quadratic, it is not surprising that the main locations of error within the support of the object are near edges. This effect of edges on noise properties has also been seen for 2D fan-beam CT in [30], which postulates that the source of the high variance near edges is the uncertainty of edge position. Accurate variance prediction near edges might require *a priori* knowledge of edge location.

We have investigated small cone angle geometries that allow us to approximate variance using the single-integral form (27). For a wider range of CT geometries, such as C-arm CT, where we cannot make a small cone angle approximation, we could instead use (24),

although this double-integral form would afford less acceleration than the single-integral form.

In these geometries, more careful attention must also be paid to $\mathcal{B}$, defined in (50) (in supplementary material), which represents the set of source positions such that a ray from the source through a voxel of interest is perpendicular to a particular frequency. For our specific geometries, (69) is an approximate closed form for (50). For more general geometries, it may be necessary to find $\mathcal{B}$ numerically or to find another approximation specific to the geometry.

VI. Conclusions and Future Work

We presented a method for approximating the variance of each voxel of a statistically reconstructed 3DCT image with quadratic regularization. This method has a computational cost that is smaller by several orders of magnitude compared to existing variance prediction methods for helical CT, while maintaining a reasonable error within regions of interest away from image edges. Generalizing the methods for nonquadratic regularizers is a challenging open problem [31].

There are a few opportunities for further acceleration. When using a spatially varying regularizer, the change in variance as a function of space is often more due to the change in the effective regularization parameter than a change in weighting or location parameters. It may be possible, then, to evaluate the integral in (27) on a coarser grid and then interpolate to all image voxels, with the factor of α_{eff} multiplying the integral producing much of the spatial variation. We also have not performed a complete analysis of strategies to numerically evaluate the integral in (27), although pilot tests have suggested that as few as 24 values of Φ suffices.

We are also investigating using our approximate factorization (23) of the LFR in other applications that could benefit from the acceleration it can provide. For example, the performance of an algorithmic observer for detecting the presence of a feature can be approximately expressed in the frequency domain [32, Ch. V]. With a fast prediction for this performance, enabled by the factorization, we could potentially optimize automatic tube current modulation methods for a particular feature-detection task, e.g., for low-contrast features in locally smooth backgrounds where the quadratic regularizer approximation is reasonable.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work was supported in part by NIH grants R01 HL-098686 and U01 EB018753 and equipment donations from Intel Corporation.

References

1. Thibault J-B, Sauer K, Bouman C, Hsieh J. A three-dimensional statistical approach to improved image quality for multi-slice helical CT. *Med. Phys.* 2007 Nov;34:4526–4544. [PubMed: 18072519]
2. Wang G, Vannier MW, Cheng PC. Iterative X-ray cone-beam tomography for metal artifact reduction and local region reconstruction. *Microscopy and Microanalysis*. 1999 Jan;5:58–65. [PubMed: 10227827]
3. De Man B, Nuyts J, Dupont P, Marchal G, Suetens P. An iterative maximum-likelihood polychromatic algorithm for CT. *IEEE Trans. Med. Imag.* 2001 Oct;20:999–1008.
4. Fessler JA. Mean and variance of implicitly defined biased estimators (such as penalized maximum likelihood): Applications to tomography. *IEEE Trans. Im. Proc.* 1996 Mar;5:493–506.
5. Cho JH, Fessler JA. Quadratic regularization design for 3D axial CT. *Proc. Intl. Mtg. on Fully 3D Image Recon. in Rad. and Nuc. Med.* 2013:78–81.
6. Gies M, Kalender WA, Wolf H, Suess C. Dose reduction in CT by anatomically adapted tube current modulation. I. Simulation studies. *Med. Phys.* 1999 Nov;26:2235–2247. [PubMed: 10587204]
7. Kalender WA, Wolf H, Suess C. Dose reduction in CT by anatomically adapted tube current modulation. II. Phantom measurements. *Med. Phys.* 1999 Nov;26:2248–2253. [PubMed: 10587205]
8. Fessler JA. Approximate variance images for penalized-likelihood image reconstruction. *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.* 1997; 2:949–952.
9. Qi J, Leahy RM. A theoretical study of the contrast recovery and variance of MAP reconstructions from PET data. *IEEE Trans. Med. Imag.* 1999 Apr;18:293–305.
10. Qi J, Leahy RM. Resolution and noise properties of MAP reconstruction for fully 3D PET. *IEEE Trans. Med. Imag.* 2000 May;19:493–506.
11. Zhang-O'Connor Y, Fessler JA. Fast predictions of variance images for fan-beam transmission tomography with quadratic regularization. *IEEE Trans. Med. Imag.* 2007 Mar;26:335–346.
12. Zhang-O'Connor Y, Fessler JA. Fast variance predictions for 3D cone-beam CT with quadratic regularization. *Proc. SPIE 6510 Medical Imaging 2007: Phys. Med. Im.* 2007:65105W:1–65105W:10.
13. Schmitt, S.; Fessler, JA. Fast variance computation for quadratically penalized iterative reconstruction of 3D axial CT images; *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf*; 2012. p. 3287-3292.
14. Schmitt SM, Fessler JA. Fast variance prediction for iterative reconstruction of 3D helical CT images. *Proc. Intl. Mtg. on Fully 3D Image Recon. in Rad. and Nuc. Med.* 2013:162–165.
15. Fessler JA, Rogers WL. Spatial resolution properties of penalized-likelihood image reconstruction methods: Space-invariant tomographs. *IEEE Trans. Im. Proc.* 1996 Sep;5:1346–1358.
16. Fessler, JA. Statistical image reconstruction methods for transmission tomography. In: Sonka, M.; Fitzpatrick, JM., editors. *Handbook of Medical Imaging, Volume 2. Medical Image Processing and Analysis*. Bellingham: SPIE; 2000. p. 1-70.
17. Sauer K, Bouman C. A local update strategy for iterative reconstruction from projections. *IEEE Trans. Sig. Proc.* 1993 Feb;41:534–548.
18. Fessler JA, Booth SD. Conjugate-gradient preconditioning methods for shift-variant PET image reconstruction. *IEEE Trans. Im. Proc.* 1999 May;8:688–699.
19. Stayman JW, Fessler JA. Regularization for uniform spatial resolution properties in penalized-likelihood image reconstruction. *IEEE Trans. Med. Imag.* 2000 Jun;19:601–615.
20. Matej S, Lewitt RM. Practical considerations for 3-D image reconstruction using spherically symmetric volume elements. *IEEE Trans. Med. Imag.* 1996 Feb;15:68–78.
21. Weideman JAC. Numerical integration of periodic functions: A few examples. *The American Mathematical Monthly*. 2002; 109(1):21–36.
22. Cho JH, Fessler JA. Regularization designs for uniform spatial resolution and noise properties in statistical image reconstruction for 3D X-ray CT. *IEEE Trans. Med. Imag.* 2015 Feb;34:678–689.
23. Gregor J. Data-driven problem reduction for image reconstruction from projections using gift wrapping. *IEEE Trans. Nuc. Sci.* 2011 Jun;58:724–729.

24. Segars WP, Mahesh M, Beck TJ, Frey EC, Tsui BMW. Realistic CT simulation using the 4D XCAT phantom. *Med. Phys.* 2008 Aug.35:3800–3808. [PubMed: 18777939]
25. Erdogan H, Fessler JA. Ordered subsets algorithms for transmission tomography. *Phys. Med. Biol.* 1999 Nov.44:2835–2851. [PubMed: 10588288]
26. Kim D, Pal D, Thibault J-B, Fessler JA. Improved ordered subsets algorithm for 3D X-ray CT image reconstruction. *Proc. 2nd Intl. Mtg. on image formation in X-ray CT.* 2012:378–381.
27. Zerhouni EA, Boukadoum M, Siddiky MA, Newbold JM, Stone DC, Shirey MP, Spivey JF, Hesselman CW, Leo FP, Stitik FP. A standard phantom for quantitative CT analysis of pulmonary nodules. *Radiology.* 1983; 149(3):767–773. [PubMed: 6647853]
28. Goodsitt MM, Chan H-P, Way TW, Larson SC, Christodoulou EG, Kim J. Accuracy of the CT numbers of simulated lung nodules imaged with multi-detector CT scanners. *Med. Phys.* 2006 Aug.33:3006–3017. [PubMed: 16964879]
29. Chang Z, Zhang R, Thibault J-B, Sauer K, Bouman C. Statistical x-ray computed tomography from photon-starved measurements. *Proc. SPIE 9020 Computational Imaging XII.* 2014:90200G.
30. Koehler T, Proksa R. Noise properties of maximum likelihood reconstruction with edge-preserving regularization in transmission tomography. *Proc. Intl. Mtg. on Fully 3D Image Recon. in Rad. and Nuc. Med.* 2009:263–266.
31. Ahn S, Leahy RM. Analysis of resolution and noise properties of nonquadratically regularized image reconstruction methods for PET. *IEEE Trans. Med. Imag.* 2008 Mar.27:413–424.
32. Schmitt, S. PhD thesis. Ann Arbor, MI, 48109-2122, Ann Arbor, MI: Univ. of Michigan; 2015. Fast variance prediction for iteratively reconstructed CT with applications to tube current modulation.

Fig. 1.

Three slices of standard deviation maps for simulated reconstructions using spatially varying regularization; the red line in (c) and (d) indicates the profile used in (e). Scales in modified Hounsfield units (air is zero, water is 1000). In panels (a)–(d), the upper left is axial, the lower left is coronal, and the upper right is sagittal. Subsequent figures use the same convention.

Fig. 2.

Three slices of standard deviation maps for simulated reconstructions using uniform regularization; the red line in (c) and (d) indicates the profile used in (e). Scales in modified Hounsfield units.

Fig. 3.

Three slices of standard deviation maps for real reconstructions using spatially varying, quadratic regularization (Hounsfield units). Coronal and sagittal slices were stretched in the trans-axial direction by a factor of two for visualization. In (c), the scale is modified HU. In (e), the thin black line indicates empirical SD; the blue line indicates the DFT-based prediction; the thick black line indicates our predicted SD.

Fig. 4.

Three slices of standard deviation maps for real reconstructions using spatially uniform, Huber-penalized regularization (Hounsfield units). Coronal and sagittal slices were stretched in the trans-axial direction by a factor of two for visualization. In (c), the scale is modified HU. In (e), the thin black line indicates empirical SD; the blue line indicates the DFT-based prediction; the thick black line indicates our predicted SD. Dashed red lines indicate locations of edges between voxels larger than 10 HU.

Fig. 5.
Cumulative distribution of error of predicted standard deviation versus empirical standard deviation for simulated reconstructions in Figs. 1–2.

TABLE I

Computation time of variance prediction methods (CPU seconds); both cases used the spatially varying regularization in Sect. III-C.

	Empirical	DFT-based	Proposed
Simulation	$1.64 \cdot 10^7$	$7.23 \cdot 10^8$	$1.21 \cdot 10^3$
	(111 realizations)		
Real data	$3.63 \cdot 10^5$	$1.07 \cdot 10^8$	$6.73 \cdot 10^2$
	(10 scans)		